REMARKS

Claims 1-15 are currently active.

Claim 8 has been amended. Antecedent support for the amendment to Claim 8 is found in Claim 1.

Claim 15 has been added.

The Examiner has rejected Claims 8 and 9 as being anticipated by Lee. Claim 8 has been amended to include the limitation of "attempting to form a connection in an ATM network satisfying original quality of service requirements of a UPC associated with the circuits". Lee does not teach or suggest anything all in regard to UPC. Accordingly, Claim 8, and Claim 9 which is dependent to parent Claim 8, are not anticipated by Lee.

The Examiner has rejected Claims 1-6 and 9-11 as being unpatentable over Lee in view of Dighe. Applicants respectfully traverse this rejection.

Referring to Lee, there is disclosed a method for dynamic connection management in integrated communication networks. Lee teaches in a connection-oriented

communication network, the transfer information between two end-users is accomplished by network functions that select and allocate the resources along an acceptable path. The logical association between the communicating end-users is referred to as a call. The chain of associated network resources that support a call is referred to as a connection. Connection management is a network function that is responsible for setting up, maintaining, and taking down connections. Each call request is issued with a set of quality of service requirements that govern the resource allocation for the desired connection. Conventional allocation of network resources for connection management is by and large static. See column 1, lines 15-29. Static allocation of network resources is inefficient, and is inadequate for a dynamic networking environment where the user requirements and the quality characteristics of network resources are not static. See column 1, lines 4-45. Lee teaches that there is a need for a system and method that addresses the problem of managing connections, and a connection-oriented integrated communications network, to maintain acceptable quality of service for the entire duration of each connection, to make quality of service improvements whenever it is possible, as well as to recover from connection failures and pre-emption. See column 2, lines 51-58.

Lee teaches the quality of service requirements for a connection specified in terms of performance constraints, resource constraints and priority constraints. A performance constraint may be non-negotiable or negotiable. A non-negotiable performance

constraint is specified with an acceptable value that separates the entire range of possible performance values into two contiguous regions: unacceptable region, an acceptable region. A negotiable performance constraint is specified with a range of values bounded between an acceptable value and requested value. The requested value divide the acceptable region into two regions. The region bounded by the acceptable value together with the requested value is referred to as the agreeable region. The other region is referred to as the overkill region. During call set up, routing computes a path, and derives from an available value for each performance parameter. If the available value falls within the unacceptable region, the call is rejected. For non-negotiable performance constraints, if the available value falls within the acceptable region, the susceptible value is taken to be the agreed value. During a negotiation, if the available value falls within the over kill region, a call is accepted. But only the requested value would be offered as the real value. See column 3, lines 43-column 4, line 20.

Resource constraints are specified with multi-level preferences but translated into acceptable and requested resource constraint sets. Each resource constraints set contains a subset of all resource attributes values. Resource constraints may also be directly specified in terms of these two constraint sets. During call set up, a routing first tries to define a path that satisfies the requested resource constraints set and other quality service constraints. If this attempt fails, the routing will use the acceptable resource constraint set for subsequent fall back path, and computations. When a path is accepted, the resource constraints set used in the

final path computation is referred to as the effective resource constraint set. See column 4, lines 35-51.

Priority constraints are specified in terms of three connection priorities: establishment of priority, retention priority, and re-establish priority. Each priority may assume one of a predetermined number of levels during call. During call setup, pre-emption is permitted only after routing has tried in vain to find an acceptable path without resorting to pre-emption. See column 4, lines 52-58.

Lee teaches a novel approach to connection management by providing an enhancement procedure coupled with a free optimization procedure, update rules for negotiable performance constraints, and extension of the pre-emption mechanism. See column 4, lines 59-64. It is within this context, that the teachings of Lee must be considered.

Lee teaches when a call arrives it enters the establishment state. If there are not enough resources to support the call, the call will be rejected. Upon successful connection establishment, the call enters the information transfer state. When a call in the information transfer state is completed, it enters the release date, and the connection is subsequently taken down. When a call is in the information transfer state, its agreed performance values may fall out of match due to facility changes, link failures, and the like. As long as the quality of

service degradation is within acceptable limits, the call remains in the information transfer state. If any degradation leads to a connection failure, or the call is pre-empted, the call enters the re-establishment state, so that the network may attempt to reestablish the connection. When a call is in the reestablishment state, the network attempts to find a new acceptable path for connection reestablishment. Upon successful reestablishment, the call reenters the information transfer state. The length of time in which a connection reestablishment attempt may be repeated is limited by a connection reestablishment delay. Beyond this delay, the reestablishment procedure is aborted, and the call enters the release date.

A call in the information transfer state is subject to enhancement to restore degraded performance values back to their agreed levels, and to avoid the use of resources that cannot satisfy the requested resource constraints. In the process of enhancement, the connection remains intact until a desired path is found. Enhancement involves finding a new path for the connection such that at least one agreed value is renegotiated to a more desirable value and none of the other agreed values is degraded. In addition, the requested resource constraints are also satisfied. See column 5, lines 5-58.

Lee teaches the enhancement procedure is coupled to a reoptimization procedure. Reoptimization is permitted only when the correctly agreed quality of service for each performance constraint is maintained at the most desirable level possible, subject to the

appropriate update rules for the performance constraint. When a call is pre-empted, it enters the re-establishment state for reroute. There are three possible outcomes for connection reestablishment: hard pre-emption, soft pre-emption and disguise pre-emption. The reestablishment may fail, whereupon the call enters the release state and is subsequently disconnected. This outcome is referred to as a hard pre-emption. Should the pre-empted call be successfully reestablished, there are either not enough resources for the pre-empted call to be reestablished with its previously agreed throughput, or there are not enough resources such that its previously agreed throughput may be compromised. In the latter case, the reestablished call is said to have suffered soft pre-emption to make room for a higher priority call. In the former case, since the connection has suffered but a minor glitch while it is being rerouted, this event is referred to as a disguised pre-emption. In both cases, the call returns to the information transfer state. Lee teaches for negotiable performance constraints, the susceptible and requested values may be updated for reestablishment, enhancement and reoptimization, according to one of the following rules: continuation rule, relaxation rule, and enhancement rule. In accordance with the continuation rule, both the requested and acceptable values associated with a given negotiable performance constraint are set to the previously agreed value, so that this value continues to be guaranteed. In accordance with the relaxation rules, the requested value is replaced by the previously agreed value. In accordance with the enhancement role, the susceptible value is replaced by the previous agreed value. See column 6, lines 18-65.

Referring to Dighe, there is disclosed a UPC-based traffic control framework for ATM networks. Dighe teaches traffic classification is based on traffic characteristics of the offered load and quality of service requirements. The framework is based on the assumption that traffic management on a per-class basis is more realistic than traffic management on a per connection or per virtual channel basis. It is assumed that the service provider will define a set of m classes of service. Each class of service is targeted towards a set of specific applications, whose traffic characteristics and quality of service requirements from the network are similar. Associated with each class is a traffic parameter vector. The traffic parameter represents the traffic characteristics of sources that subscribe to a given service class. Elements of the traffic parameter vector will be quantities such as the peak rate or burst size. Preferably, the traffic parameter vector is assumed to be the same as the traffic shaping vector defined by the ATM Forum. This is based on the dual-leaky bucket filter. See column 3, lines 12-43. It is within this context that all the teachings of Dighe are based on the dual-leaky bucket filter that the teachings of Dighe must be considered.

Dighe teaches a dual leaky bucket is define by three parameters namely: the peak rate, the sustained rate, and the compliance burst size. Once a source declares these parameters, the dual leaky bucket implementation of the UPC function at the UNI ensures that the source transmits no more than the compliant burst size of cells at the peak rate. If the burst size is larger than the compliant burst size, then the first compliant burst size cells are

transmitted at the peak rate. The remaining cells in the burst are only transmitted at the sustained rate which can be at most equal to the peak rate. Hence, once the three parameters of the UPC function are known, the switch can determine a worst-case traffic profile for the offered traffic. Based on this worst-case traffic profile, the switch can decide to either accept the call, reject the call, or accept a call after suitably modifying the declared UPC parameters. See column 5, line 55-column 6, line 5. The remaining description found in Dighe focuses on the three parameters of the dual leaky bucket in regard to the specific architecture and equipment described by Dighe.

Applicant is pointing out this dual leaky bucket basis of the teachings of Dighe because of a criticalness of the system to be able to perform the dual leaky bucket implementation of the UPC function to be able to practice the teachings of Dighe.

The Examiner has concluded that the combination of Lee and Dighe arrive at applicants' claimed invention. Applicants respectfully submit that the teachings of Lee and Dighe cannot be combined because not only is there no teaching in the references themselves to combine them, but the teachings cannot be taken out of the context in which they are found, and by doing so, would be unworkable.

It is black letter patent law that there must be some teaching or suggestion in the references themselves to combine the references, and here, there is no such teaching or suggestion. Nowhere does Lee even consider the use of the UPC function, and this follows, because Lee focuses on a novel approach to connection management by providing an enhancement procedure coupled with a reoptimization procedure, update rules for negotiable performance constraints, an extension of the pre-emption mechanism. See column 4, lines 59-64. There is no need for UPC in Lee and it is not considered. In fact, there is no reason why one skilled in the art would even apply the UPC function (if they even could) to the teachings of Lee. Similarly, the dual leaky bucket that is the basis of the UPC function taught by Dighe has no need for the novel approach to connection management involving an enhancement procedure coupled with a reoptimization procedure, update rules for negotiable performance constraints, and extension of the pre-emption mechanism, as taught by Lee. This also follows, since Dighe is not interested in any such enhancement procedure coupled with the reoptimization procedure and has no need for the same, since there is no place for the same in the operation of the system taught by Dighe and the dual leaky bucket. Applicants submit that the only reason that the Examiner is combining the teachings of Lee and Dighe is from the hindsight of applicants' claimed invention, and this is contrary to patent law. The Examiner cannot use applicants' claims as a road map to find the elements and limitations of applicants' claims in various references, and having found them, concluding that applicants' claimed invention is arrived at.

Furthermore, not only is there no teaching to combine Lee and Dighe, but when teachings from different references are combined, they must be done so in the context in which they are found. Applicants have just explained above the specific context of each of the references. The context of Lee is in regard to a novel approach to connection management providing an enhancement procedure coupled with the reoptimization procedure, and all the associated teachings and operation of Lee follows. The context of Dighe involves the UPC function that is based on the dual leaky bucket. These contexts have nothing to do with each other. The parameters that are required for the dual leaky bucket, and are identified by Dighe cannot be found and are not identified at all in the teachings of Lee. To combine the teachings of the UPC that the Examiner suggests is found in Dighe, requires how the UPC is performed in Dighe must also be applied under patent law to the teachings of Lee. By doing so, not only does it not make any sense to apply the dual leaky bucket somehow or other to the operation of Lee, but it would require Lee to be completely redesigned and require significant amount of experimentation and testing to even possibly be workable. This only further supports the non obviousness of applicants' claimed invention. Accordingly for this reason, applicants' claimed invention is not obvious from, and is patentable over Lee in view of Dighe.

Moreover, precisely because of how different the contexts of each of these references are that one relies completely on the dual leaky bucket, and the other one has nothing at all to do with the dual leaky bucket but applies a totally different approach to

connection management, one skilled in the art cannot combine these references and make the result operational.

Accordingly, Claims 1-6 and 9-11 are patentable over Lee in view of Dighe.

The Examiner has rejected Claim 7 as being an patentable over Lee in view of Dighe and further in view of Burns. Applicants respectfully traverse this rejection. Burns adds nothing in relevant part to the teachings of Lee or Dighe to arrive at applicants' claimed invention. The applied art of record does not teach or the suggest "a controller which attempts to establish the circuit according to original quality of service requirements of the UPC associated with the circuit". Accordingly, Claim 7 is patentable over the applied art of record.

The Examiner has rejected Claims 10-14 as being unpatentable over Lee. For the reasons explained above, Lee does not teach or suggest anything at all in regard to the UPC function. Accordingly, Claims 10-14, which are dependent to parent Claim 8, are patentable over Lee.

In view of the foregoing amendments and remarks, it is respectfully requested that the outstanding rejections and objections to this application be reconsidered and withdrawn, and Claims 1-15, now in this application be allowed.

CERTIFICATE OF MAILING

I hereby certify that the correspondence is being deposited with the United States Postal Service as first class mail in an envelope addressed to: Assistant Commissioner for Patents, Washington, D.C. 20231, on 7/17/03

Ansel M. Schwartz
Registration No. 30,587
7/17/03
Date

Respectfully submitted,

MARK HELD, ET AL.

By Ansel Schwartz

Ansel M. Schwartz, Esquire
Reg. No. 30,587
One Sterling Plaza
201 N. Craig Street
Suite 304
Pittsburgh, PA 15213
(412) 621-9222

Attorney for Applicants

<u>**Version with markings to show changes made to the claims**</u>

8. A method for establishing circuits in an ATM network comprising the steps of:

attempting to form a connection in an ATM network satisfying original quality of service requirements <u>of a UPC associated with the circuits</u>;

rejecting the formation of the circuit due to resources of the ATM network not being available to meet the original quality of service requirements of the circuit;

relaxing automatically the quality of service requirements of the circuit; and

creating the circuit in the ATM network subject to the relaxed quality of service requirements.